

Overlay Routing for Fast Video Transfers in CDN

Paolo Medagliani*, Stefano Paris*, Jérémie Leguay*, Lorenzo Maggi*, Xue Chuangsong†, Haojun Zhou†

*Mathematical and Algorithmic Sciences Lab, France Research Center, Huawei Technologies Co. Ltd.

20 Quai du Point du Jour, 92100 Boulogne-Billancourt, France

†Carrier Software Unit, Huawei Technologies Co. Ltd., Nanjing, China

Abstract—Content Delivery Networks (CDN) are witnessing the outburst of video streaming (e.g., personal live streaming or Video-on-Demand) where the video content, produced or accessed by mobile phones, must be quickly transferred from a point to another of the network. Whenever a user requests a video not directly available at the edge server, the CDN network must 1) identify the best location in the network where the content is stored, 2) set up a connection and 3) deliver the video as quickly as possible. For this reason, existing CDNs are adopting an overlay structure to reduce latency, leveraging the flexibility introduced by the Software Defined Networking (SDN) paradigm. In order to guarantee a satisfactory Quality of Experience (QoE) to users, the connection must respect several Quality of Service (QoS) constraints. In this paper, we focus on the sub-problem 2), by presenting an approach to efficiently compute and maintain paths in the overlay network. Our approach allows to speed up the transfer of video segments by finding minimum delay overlay paths under constraints on hop count, jitter, packet loss and relay node capacity. The proposed algorithm provides a near-optimal solution, while drastically reducing the execution time. We show on traces collected in a real CDN that our solution allows to maximize the number of fast video transfers.

I. INTRODUCTION

Internet video traffic is expected to grow fourfold in the next four years up to represent 82 percent of all IP traffic (business and consumer) by 2020 [1]. While a large variety of Video on Demand (VoD) and video-streaming services has emerged in the past years, the field continues to evolve rapidly. The ways people are watching video is constantly evolving and it is driven by mobile usage. For instance, live streaming embedded in social media platforms is a relatively new phenomenon, but this technology is finding more and more support with services such as Facebook Live or Periscope.

With the explosion of streaming services that deliver Internet video to the TV and other device endpoints, Content Delivery Networks (CDN) have prevailed as the dominant technology to deliver such content. Globally, 72 percent of Internet video traffic will cross CDN by 2019. The largest over-the-top player Akamai currently has over 170,000 edge servers located in over 1300 networks in 102 countries [2]. At a smaller scale, Internet Service Providers are also deploying their own infrastructure, referred to as Telco CDN, as an evolution of IPTV and VoD systems. While CDNs have been traditionally deployed for the distribution of static content at large scale, they are evolving to support a large variety of content types like web applications, teleconferencing and live

video streaming which cannot always be cached due to their dynamic nature.

In such dynamic and large-scale settings, delivering content over the Internet with latency and reliability constraints becomes a real challenge. Indeed, the Internet is best-effort with routing policies that do not address fined-grained needs of applications and that are often guided by business relationships on large traffic volumes. The Triangle Inequality Violation (TIV) [3] is a well known consequence of such policies. The minimum delay path is almost never the one established by the underlying routing system. In addition, outages are happening all the time in the Internet due to cable cuts, mis-configured routers, DDoS attacks, power outages, or natural disasters [4]. Even if the Internet becomes flatter [5] with content service providers buying direct connectivity closer to their end users, CDN operators are still fighting against TIV and best effort routing policies. In reaction, overlay networks, such as RON [6], have been introduced to provide low latency and reliable connectivity over the Internet. Similarly, CDN operators deploy a three-tier architecture composed of *origin servers* that create the content, *edge servers*, which clients access to consume the content, and an *overlay network* that transports the content from the origins to the edges.

CDN solutions are composed of building blocks such as a caching system to store the most popular contents at the network edge, a load balancing system, integrated within a Domain Name System (DNS) server, to redirect client requests to the closest edge server and an overlay routing system to transport content at low latency and high reliability. Following the ongoing transformation of network architectures with Software Defined Networks (SDN) [7], CDN are adopting flow-oriented and centralized controllers [8] to manage video traffic especially. The (logically) centralized control aims at improving the Quality of Experience (QoE) perceived by end users. The challenge for such an overlay network controller is to quickly find paths in the overlay when new demands arrive and to maintain a good routing configuration over time so that the transfer time of video segments is minimized.

This paper presents an efficient algorithm to maintain minimum delay overlay paths with multiple QoS constraints on capacity, jitter and packet loss. These optimization criteria have been carefully selected to speed-up the transfer of video segments, knowing that TCP or QUIC [9] is used under HLS or MPEG-DASH [10] for live and regular streaming. Our algorithm is based on the decomposition of the original problem

into two simpler problems, namely finding a minimum delay path that satisfies all QoS constraints for a single demand (QoS path computation), and computing the best set of paths for all demands (constrained multi-commodity flow). We present an evaluation of this algorithm over measurements collected in a real Telco CDN.

The rest of the paper is structured as follows. Sec. II provides an overview of the related work. Sec. III introduces the system model considered and formulates the problem as an MILP. Sec. IV describes the algorithms that we propose to solve the admission and routing maintenance problems. Sec. V illustrates the numerical evaluation of the proposed algorithms. Finally, concluding remarks are provided in Sec. VI.

II. RELATED WORKS

Overlay routing has received a lot of attention, especially in the domains of peer-to-peer, conferencing and CDNs.

QoS routing for multi-layer video streams has been proposed to compute disjoint overlay paths for each sub-stream to increase reliability [11]. Thi et al. [12] allocate all sub-streams at once by minimizing an estimate of the end-to-end video quality distortion. They minimize the end-to-end probability of video stall based on packet loss and latency.

Distributed path computation algorithms have been proposed for peer-to-peer conferencing applications to build and maintain application layer multicast trees [13]. They gradually maintain a tree for each multicast session by defining join and leave procedures. Conversely, our work suits for the CDN case with a centralized controller platform [8].

Andreev et al. [14] introduce an overlay network for live streaming with edge servers, reflectors and source nodes. They propose a linear relaxation and rounding based algorithm to select reflectors. As reflectors can split multimedia streams to serve multiple receivers, the algorithm build an overlay forest to connect sources to receivers. To increase reliability, they also ensure that each source is served by two reflectors. Their primary objective is to optimize the cost of the infrastructure. Similarly, Zhou et al. [15] proposed an algorithm to find capacity and delay bounded minimum cost overlay forests.

Our work differs from state of art by focusing on delay minimization subject to constraints on jitter and packet loss. In addition, it considers that edge servers act at the same time as source, relay and end points in the overlay network. In this context, we provide an algorithmic framework for the online and global control of fast video transfers in CDN.

III. PROBLEM FORMULATION

This section introduces the system model and the path computation problem.

A. System model

A typical CDN framework is depicted in Fig. 1. We consider the presence of several content sources that stream videos, both live and on demand, to connected remote end users using the existing overlay network. Instead of connecting directly the end users to the source, end users connect to an edge

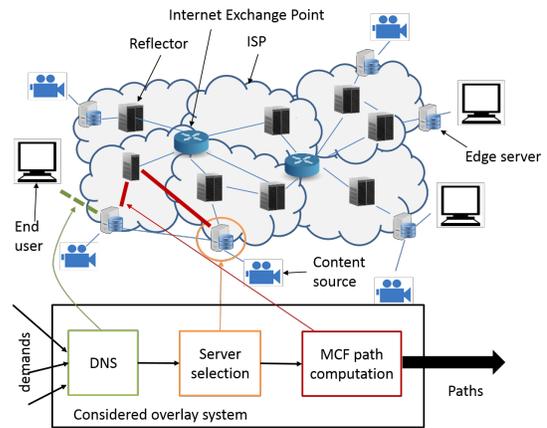


Figure 1: System model of a CDN overlay framework.

server. In fact, since the edge server is able to replicate the same stream towards different end users, this results in a load reduction at sources and a better bandwidth utilization in the overlay. The choice of the best edge server which an end user must be connected to is typically handled by a DNS or HTTP proxy system. By leveraging different information such as geographical locations, content availability and edge server load, it redirects end user connection requests towards the most suitable edge server. Once the best server has been identified, two cases arise: (i) the edge server has already the content available and it can serve the end user, (ii) the content is not available at the edge server. In the latter case, the edge server needs to retrieve the requested video from either the origin server or another edge server that already has the content available. To this end, it is necessary to compute the best overlay path to transfer the content to the edge server querying for it. Fig. 1 shows a schematic description of the interactions among the different overlay components.

While existing overlays are composed of edge servers belonging to the same provider, we consider also the case where they can be extended with nodes placed at Internet eXchange Points (IXP) thanks the emerging technology of Software-Defined Exchange (SDX). Fressancourt et al. [16] have shown that IXP can be used as third-party routing inflection points to enhance an existing overlay network. This way, CDN overlays with even a few internal nodes can reach a high level of path diversity using external relays.

The system model presented above applies to two use cases: (i) video on demand (VoD) and (ii) personal live streaming (PLS). In the former scenario, end users request a video from a content provider. The goal of the system is to guarantee that the user is served as quickly as possible. If the content is already available at the edge server the problem is trivial and it only concerns the connection between end user and edge server. Instead, if the content is not directly available at the edge server, then the problem is equivalent to computing a path which minimizes the latency between two edge servers in the network and which respects some QoS constraints. The video content is either retrieved from the source or from the

Symbol	Description
\mathcal{N}	Nodes (network devices).
\mathcal{E}	Edges (network links).
\mathcal{K}	Set of demands (i.e., commodities).
r^k	Transmission rate for demand $k \in \mathcal{K}$.
$d_{i,j}$	Delay of edge $(i,j) \in \mathcal{E}$.
$b_{i,j}$	Capacity of edge $(i,j) \in \mathcal{E}$.
$f_{i,j}$	Prob. of successful transmission for edge (i,j) .
F^k	Minimum prob. of successful transmission for k .
$z_{i,j}$	Jitter of edge $(i,j) \in \mathcal{E}$.
Z^k	Maximum jitter for demand $k \in \mathcal{K}$.
N_i	Maximum processing rate for node $i \in \mathcal{N}$.

Table I: Notations for input parameters.

cache of another edge server.

In the PLS scenario, instead, content is generated and streamed by users, as it may happen for instance with Facebook Live. Other end users willing to watch the content connect to the overlay network in order to retrieve it. At this point, a similar situation to VoD arises. The overlay network first identifies the best edge server which the user must connect to; if the edge server does not have the content available, then it requests it from either the origin server or from another edge server [4]. The system computes and maintains the fastest path between the content source and the demanding edge server. Throughout this paper, we consider that an external element has already chosen the best edge and origin servers from which the content must be retrieved. Hence, we will only focus on the path computation problem.

In both cases, videos are streamed using HTTP Live Streaming or MPEG-DASH over TCP or QUIC. Minimizing transfer durations then translates into maximizing the throughput of each transport session. For TCP, the throughput can be approximated by the following formula $\frac{MSS \cdot C}{RTT \cdot \sqrt{p}}$ [17] which takes in to account the Maximum Segment Size (MSS), the Round-Trip Time (RTT) and the packet loss probability p . As a consequence, our path finding and maintenance algorithm aims at bounding packet loss and jitter while minimizing RTT. The parameters are continuously monitored by an active monitoring system. In addition, as the relaying of flows induces a burden on overlay nodes, we also consider a maximum amount of traffic that each one can process.

We point out that the demands accepted by our algorithm will be routed in the CDN overlay, following the computed paths. For the refused demands, instead, they will be accepted anyway by the CDN overlay but using the direct path between origin and edge server in a best-effort way.

B. Mathematical formulation

In this paper we model the CDN overlay network as a weighted directed graph $G = (\mathcal{N}, \mathcal{E})$, where \mathcal{N} is the set of nodes and \mathcal{E} denotes the set of edges. Each directed edge $(i,j) \in \mathcal{E}$ is characterized by its capacity $b_{i,j}$, delay $d_{i,j}$, jitter $z_{i,j}$ and successful packet transmission probability $f_{i,j}$. Each node $i \in \mathcal{N}$ has a maximum processing rate N_i .

We consider a set \mathcal{K} of video-streaming connection demands which need to be routed in the network. Demand k is identified

by a source node $s^k \in \mathcal{N}$, destination node $t^k \in \mathcal{N}$ and transmission rate r^k . Table I summarizes the notation used throughout the paper. Our primary objective is to accept the maximum number of demands into the system; our secondary goal is to minimize the total delay, while each demand must fulfill all hard constraints on link and node capacity, jitter and packet loss probability. Moreover, at most one reflector can be used for each demand. This translates into a Multi-Commodity Flow (MCF) problem, that can be expressed either via a link- or path-based formulation, according to the needs. We start off with the link-based formulation, that sets the decision variable $x_{i,j}^k = 1$ if and only if the directed edge (i,j) is used for routing demand k , i.e.,

$$\min_x \sum_{k \in \mathcal{K}} \sum_{(i,j) \in \mathcal{E}} x_{i,j}^k d_{i,j} + M \sum_{k \in \mathcal{K}} \sum_{(i,j) \in \bar{\mathcal{E}}} x_{i,j}^k \quad (1)$$

$$\text{s.t.} \quad \sum_{j:(i,j) \in \mathcal{E} \cup \bar{\mathcal{E}}} x_{i,j}^k - \sum_{j:(j,i) \in \mathcal{E} \cup \bar{\mathcal{E}}} x_{j,i}^k = \gamma_i^k, \quad \forall i \in \mathcal{N}, k \in \mathcal{K} \quad (2)$$

$$\sum_{k \in \mathcal{K}} x_{i,j}^k r^k \leq b_{i,j}, \quad \forall (i,j) \in \mathcal{E} \cup \bar{\mathcal{E}} \quad (3)$$

$$\sum_{(i,j) \in \mathcal{E} \cup \bar{\mathcal{E}}} x_{i,j}^k z_{i,j} \leq Z^k, \quad \forall k \in \mathcal{K} \quad (4)$$

$$\sum_{k \in \mathcal{K}} \sum_{j:(i,j) \in \mathcal{E} \cup \bar{\mathcal{E}}} x_{i,j}^k r^k \leq N_i, \quad \forall i \in \mathcal{N} \quad (5)$$

$$\sum_{k \in \mathcal{K}} \sum_{j:(j,i) \in \mathcal{E} \cup \bar{\mathcal{E}}} x_{j,i}^k r^k \leq N_i, \quad \forall i \in \mathcal{N} \quad (6)$$

$$\sum_{(i,j) \in \mathcal{E} \cup \bar{\mathcal{E}}} x_{i,j}^k \log f_{i,j} \geq \log F^k, \quad \forall k \in \mathcal{K} \quad (7)$$

$$\sum_{(i,j) \in \mathcal{E} \cup \bar{\mathcal{E}}} x_{i,j}^k \leq 2, \quad \forall k \in \mathcal{K} \quad (8)$$

$$x_{i,j}^k \in \{0, 1\}, \quad \forall i, j, k. \quad (9)$$

We remark that, via the formulation in (1-9), we implicitly augmented the original graph G into a clique, where all the new (artificial) edges $\bar{\mathcal{E}}$ have a large delay M . More specifically, M should be set as larger than $2|\mathcal{K}| \max_{i,j} d_{i,j}$ in order to prioritize the maximization of the number of accepted demands over the delay minimization goal. Eq. (2) describes the standard flow conservation constraints, where $\gamma_i^k = 1$ if $i = s^k$, $\gamma_i^k = -1$ if $i = t^k$ and $\gamma_i^k = 0$ otherwise. Eqs. (3) and (4) account for the capacity and jitter hard constraints, respectively. Z^k is the maximum jitter value for demand k . Expressions (5),(6) ensure that each node i processes traffic at a rate not exceeding N_i . Eq. (7) claims that the probability of successful transmission for demand k is at least F^k . Finally, Eq. (8) translates the requirement that at most one reflector is used for each demand (i.e., max. number of hops equals 2).

We now present the path-based formulation for our optimization problem, that is equivalent to the link-based one in Eqs. (1)-(9). As explained in the next section, this formulation enables the decomposition of the original overlay routing problem into simpler and smaller sub-problems that can be solved more efficiently. We define \mathcal{P}^k as the set of paths from source s^k to destination t^k fulfilling the constraints on number

of hops (≤ 2), jitter ($\leq Z^k$) and probability of successful transmission ($\geq F^k$). We call $\mathcal{P}_i^k \subseteq \mathcal{P}^k$ the set of paths visiting node $i \in \mathcal{N}$. Similarly, $\mathcal{P}_e^k \subseteq \mathcal{P}^k$ is the set of paths crossing edge $e \in \mathcal{E}$. Moreover, let d_p be the total delay of path p , by accounting that edges not in \mathcal{E} have delay M . Then, the path-based formulation for MCF writes as follows:

$$\min_y \sum_{k \in \mathcal{K}} \sum_{p \in \mathcal{P}^k} y_p d_p \quad (10)$$

$$\text{s.t.} \sum_{k \in \mathcal{K}} \sum_{p \in \mathcal{P}_e^k} y_p r^k \leq b_e, \quad \forall e \in \mathcal{E} \quad (11)$$

$$\sum_{k \in \mathcal{K}} \sum_{p \in \mathcal{P}_i^k} y_p r^k \leq N_i, \quad \forall i \in \mathcal{N} \quad (12)$$

$$\sum_{p \in \mathcal{P}^k} y_p = 1, \quad \forall k \in \mathcal{K} \quad (13)$$

$$y_p \in \{0, 1\} \quad (14)$$

where $y_p = 1$ whenever path $p \in \mathcal{P}^k$ is used for $k \in \mathcal{K}$.

We observe that the two set of constraints (11)-(12) represent the transmission and processing capacity limits of links and nodes, respectively.

IV. PATH FINDING ALGORITHM FOR FAST TRANSFERS

The overlay routing problem is NP-Hard, since it generalizes the Multi-Commodity Integral Flow problem, which is known to be NP-hard. Therefore, the computational time steeply increases with the network size and the number of demands. To solve the overlay routing problem even in large scale scenarios, we propose to decompose the original problem into simpler subproblems that can be solved more efficiently. More specifically, we first relax the integrality constraints for the variable y_p , enabling the use of multiple paths for the routing of each demand, and then we design a column generation algorithm [18] to solve the underlying MCF problem. In order to deal with the QoS constraints, we use a pseudo-polynomial algorithm to generate only shortest paths that satisfy constraints (4), (7), and (8). Finally, we design a randomized method to assign a single path to those demands whose LP solution consists in splitting the traffic over multiple paths.

A. Solving the Constrained MCF problem

The column generation technique enables us to consider only a subset of decision variables in the primal formulation at each iteration, and it uses the dual formulation to include only those variables that can improve the objective function. Indeed, from the duality theory we know that every feasible point to the dual problem μ^* gives a lower bound on the optimum value of the primal y^* , and every feasible point to the primal problem y^* gives an upper bound on the optimal value of the dual μ^* . Therefore, a feasible primal solution y^* is optimal if the corresponding dual point μ^* is feasible. Our algorithm exploits this property in order to consider, for each demand, only a small set of variables representing feasible paths and add new variables to the primal formulation as long as the corresponding dual solutions are unfeasible. In our problem,

the vector of dual variables $\mu = [\lambda, \sigma]$ is split into two different sets of variables, where λ corresponds to the capacity constraints (11)-(12) and σ corresponds to constraint (13), which indicates that a subset of \mathcal{P}^k is used to route a demand.

We underline that, as long as the paths generated during the column generation procedure satisfy all QoS constraints, the final solution computed by our algorithm is optimal for the LP relaxation of the overlay routing problem, in the sense that each demand is fully satisfied by one or multiple paths. To this aim, we use the GEN-LARAC (GLC) algorithm as a subroutine for solving the constrained shortest path problem [19], since it computes in pseudo-polynomial time a path that satisfies all QoS constraints (4), (7), and (8). Therefore, at the end of the column generation algorithm we only have to choose a single path for each demand, without reconsidering all constraints of the original problem.

If $\lambda \geq 0$, the constraints of the dual problem can be formulated as follows:

$$\sigma_k^* - \sum_{e \in p} r_k \lambda_e^* \leq d_p, \quad \forall k \in \mathcal{K}, p \in \mathcal{P}_k, \quad (15)$$

Eq. (15) states that in a feasible point of the dual problem $\mu^* = [\lambda^*, \sigma^*]$, there exists no path for any demand such that $\sigma_k^* > \sum_{e \in p} r_k \lambda_e^* + d_p$, which can be rewritten as $\sigma_k^* > \sum_{e \in p} (r_k \lambda_e^* + d_e)$. In other words, there exists no path for any demand that can further improve the objective function (10). In contrast, if such a path exists, i.e., $\exists k \in \mathcal{K}, p \in \mathcal{P}_k: \sigma_k^* > r_k \sum_{e \in p} (r_k \lambda_e^* + d_e)$, then the dual point μ^* is unfeasible and the objective function of the primal problem can be further reduced by considering such a path (recall that the primal solution is an upper bound of the dual solution).

All we need to do at each iteration of the column generation algorithm is checking whether the condition $\sigma_k^* > r_k \sum_{e \in p} (r_k \lambda_e^* + d_e)$ holds for all possible paths of all demands (i.e., $\forall p \in \mathcal{P}_k$). We observe that the computation of $\sum_{e \in p} (r_k \lambda_e^* + d_e)$ can be performed efficiently using an algorithm for the constrained shortest path computation on the weighted graph $G(\mathcal{N}, \mathcal{E}, w(\cdot))$, where the link weight function $w(\cdot) : \mathcal{E} \rightarrow \mathbb{R}_{\geq 0}$ is computed as $w_e = r_k d_e + \lambda_e$.

We point out that, in order to maintain the feasibility of this routine, we also add some *dummy paths* on which demands, rejected by column generation, can be allocated.

B. Rounding procedure

To address the splitting of demands over multiple paths, which can be caused by the resolution of the LP relaxation of the problem (10)-(12), we introduce a randomized rounding phase at the end of the column generation algorithm, in order to convert the fractional solution into a feasible integer solution. Whenever multiple paths are used to route a demand from its origin to its destination, a single route is selected with probability equal to the portion of the demand allocated to each specific path. In particular, for each demand k that has been split, we consider all possible paths where $y_p > 0$ and

\mathcal{N}	11	Number of nodes
\mathcal{E}	82	Number of links
$b_{i,j}$	50	Link capacity [Gbps]
N_i	150	Node processing capacity [Mbps]
\mathcal{K}	82	Number of demands
r_k	50	Transmission rate [Mbps]
F^k	99%	Minimum prob. of successful transmission
Z^k	2	Maximum jitter for demand k [s]

Table II: Parameters for the Telco CDN scenario.

we select a path p according to the overall flow allocated to it. If the whole demand k can be transmitted at its nominal rate r_k over the selected path p without violating any capacity constraint, we keep only the path p by fixing the variable $y_p = 1$ and all other variables corresponding to alternative paths to 0. Finally, we reduce the capacity of the links that belong to p by the demand's rate r_k . The randomized routing procedure terminates either when a single path has been allocated to all demands or when the solution \mathbf{y} does not change between two consecutive iterations. In this latter case, all demands that are still split over multiple paths are rejected.

V. PERFORMANCE RESULTS

In this section, we analyze the performance our Column Generation (CG) approach based on GLC (CG-GLC). In order to show the distance of CG-GLC from the optimal solution, also called ‘‘optimality gap’’, we use the solution of the ILP presented in Section III-B as a benchmark.

Results are evaluated in a *Telco CDN* scenario, where we used traces collected on a real overlay network of a Telco CDN. We present a performance evaluation over time considering as indicators the percentage of accepted demands, the running time and the average delay. We point out that nodes can be either source/destination of a demand or reflector if the demand is not originated in or destined to the considered node. Without loss of generality, for each demand, we considered $F^k > 99\%$ and $Z^k < 2$ s.

All the tests have been executed on a machine with Intel Core i7-5600U 2.6 GHz with 8 GB of RAM. The C++ libraries used to build and update the network graph have been provided by Lemon. The ILP and the resolution of the reduced master problem have been carried out using CPLEX.

A. Real CDN Overlay Network

In this subsection, first we present the traces used for simulations and then we compare the performance of CG-GLC and ILP via numerical evaluations. In Table II, we list the main parameters considered in the CDN overlay scenario.

Transmission rate and node processing capacity are set to 50 Mbps and 150 Mbps, respectively, which leads to a scenario where nodes are fairly stressed. The link QoS metrics used for our experiments, as delay, jitter and packet loss, are described in Subsection V-A1.

1) *Dataset presentation:* We used real traces from a Telco CDN operator. Since the QoS of the considered overlay network varies over time, a monitoring system carries out periodic probing between overlay nodes. In this subsection,

we show the results of QoS metrics monitoring considering an aggregation period of 5 minutes.

In Fig. 2a, we show the pairwise distribution of the average delay. We can see that the support of delay distribution is 5-80 ms, while only few links present larger delay. The jitter distribution presented in Fig. 2b shows that the jitter is massively concentrated in the first bins. In fact, 90% of the jitter samples are smaller than 330 ms. Finally, as shown in Fig. 2c, most of the links have an average packet loss smaller than 0.3%, although 6% of the links have a packet loss of 3%. This is due to local perturbations that caused consistent packet losses.

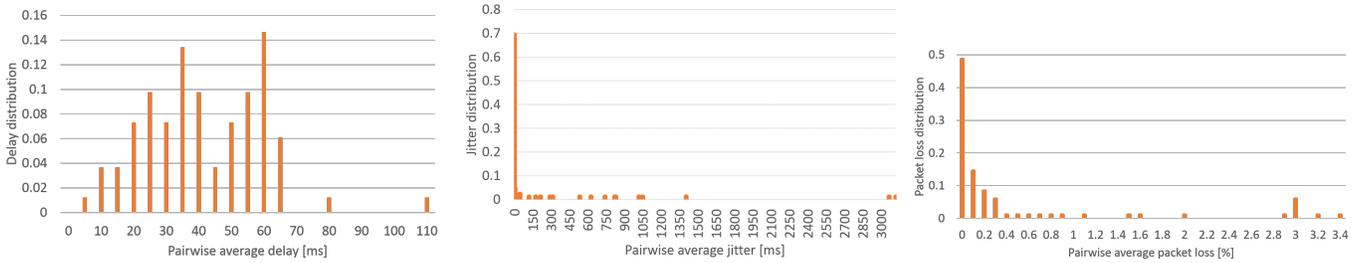
2) *Performance Evaluation:* We now compare the performance of CG-GLC and the solution of ILP in two scenarios. The former is the one described before, where an overlay network is considered (*overlay* in the Figures). The latter neglects the presence of the overlay, while only relying on the direct path between source and destination (*direct*).

We compare CG-GLC and ILP solutions in terms of percentage of accepted demands and run-time in Fig. 3a and Fig. 3b, respectively. At each time instant we try to allocate all the 82 demands. As the network capacity is not saturated, the node processing capacity is the real bottleneck of the system (i.e., constraints (5) and (6)).

We first remark that while through the overlay (*overlay*) we manage to accept in average 90% of the demands, without the overlay (*direct*) we cannot accept more than 20 % of the demands. Accepted demands are flows for which QoS constraints in term of jitter and packet loss are met. This result highlights the benefit of overlay networking compared to the case when one relies only on the underlay. Indeed, it increases path diversity and the chances to find feasible paths.

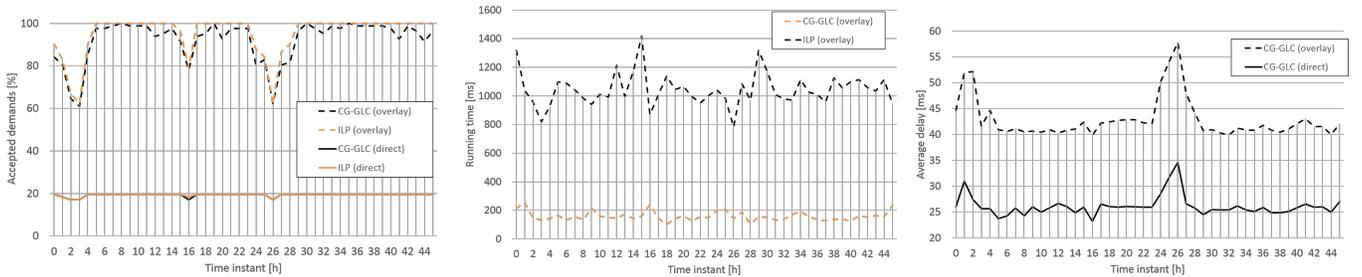
Secondly, we observe that our CG-GLC solution strikes a good performance/complexity trade-off. In fact, CG-GLC is characterized by an average optimality gap of 3% in terms of percentage of accepted demands. Moreover, CG-GLC has a running time around 10 times smaller than ILP solution, that was produced by standard commercial software CPLEX. This is due to the fact that CG-GLC, instead of solving the whole problem, focuses only on a subset of the original problem (faster to be solved), by adding to it only the solutions (i.e., paths) which improve the objective function. We point out that the CG-GLC can compute an entire network reconfiguration in less than 200 ms.

As the number of accepted demands is not the same for CG-GLC and ILP, it is not possible to carry out a fair comparison between the two approaches in terms of average delay. As shown in Fig. 3c, the average delay without overlay (*direct*) is smaller than the one with the overlay (*overlay*). This is because the number of accepted demands is smaller without the overlay and, according to the CG routine, CG-GLC tries to use first paths with low delay but, as long as new demands are accepted, they are routed on more ‘‘expensive’’ paths.



(a) Distribution of pairwise avg. delay with 5 ms bins. (b) Distribution of pairwise avg. jitter with 5 ms bins. (c) Distribution of pairwise avg. packet loss with 0.1% bins.

Figure 2: Overlay link statistics of the Telco CDN overlay.



(a) Percentage of accepted demands over time.

(b) Running time over time.

(c) Average delay over time.

Figure 3: Performance of the proposed overlay routing algorithm.

VI. CONCLUDING REMARKS

In this paper, we have addressed the problem of fast video delivery in CDN. We propose an algorithm to nearly optimally allocate and maintain paths in the overlay network. More specifically, we formulate the problem as a multi-commodity flow under several QoS constraints derived from the requirements of CDN overlays networks. The proposed solution applies to the fast video delivery problem for both personal live streaming and Video-on-Demand use cases. Our approach, based on column generation and randomized rounding, has been tested on real traces from a Telco CDN. As a benchmark, the optimal solution is computed by solving the associated ILP formulation with commercial software. Results show that our approach is an excellent compromise in terms of running time and optimality.

REFERENCES

- [1] Cisco, "Cisco Visual Networking Index: Forecast and Methodology, 2014–2019," May 2015.
- [2] B. M. Maggs and R. K. Sitaraman, "Algorithmic nuggets in content delivery," *ACM SIGCOMM CCR*, vol. 45, no. 3, pp. 52–66, 2015.
- [3] C. Lumezanu, R. Baden, N. Spring, and B. Bhattacharjee, "Triangle inequality variations in the internet," in *Proc. ACM IMC*, 2009.
- [4] E. Nygren, R. K. Sitaraman, and J. Sun, "The akamai network: a platform for high-performance internet applications," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 3, pp. 2–19, 2010.
- [5] Y.-C. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, and R. Govindan, "Are we one hop away from a better internet?" in *Proc. ACM IMC*, 2015.
- [6] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris, "Resilient overlay networks," *ACM SIGCOMM CCR*, vol. 32, no. 1, 2002.
- [7] B. A. A. Nunes, M. Mendonca, X. N. Nguyen, K. Obraczka, and T. Turletti, "A survey of software-defined networking: Past, present, and future of programmable networks," *IEEE Com. Surveys Tutorials*, 2014.
- [8] X. Liu, F. Dobrian, H. Milner, J. Jiang, V. Sekar, I. Stoica, and H. Zhang, "A case for a coordinated internet video control plane," in *Proc. ACM SIGCOMM*, 2012.
- [9] J. Iyengar, I. Swett, R. Hamilton, and A. Wilk, "QUIC: A UDP-Based Secure and Reliable Transport for HTTP/2," Internet-Draft draft-tsvwg-quic-protocol-02, Jan. 2016.
- [10] M. Seufert, S. Egger, M. Slanina, T. Zinner, T. Hoßfeld, and P. Tran-Gia, "A survey on quality of experience of http adaptive streaming," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 469–492, 2015.
- [11] A. C. Begen, Y. Altunbasak, O. Ergun, and M. H. Ammar, "Multipath selection for multiple description video streaming over overlay networks," *Sig. Proc.: Image Com.*, vol. 20, no. 1, pp. 39–60, 2005.
- [12] T. M. Thi, T. Huynh, and W.-J. Hwang, "Qos-enabled streaming of multiple description coded video over openflow-based networks," *Nonlinear Theory and Its Applications, IEICE*, vol. 6, no. 2, pp. 144–159, 2015.
- [13] M. Hosseini, D. T. Ahmed, S. Shirmohammadi, and N. D. Georganas, "A survey of application-layer multicast protocols," *Commun. Surveys Tuts.*, vol. 9, no. 3, Jul. 2007.
- [14] K. Andreev, B. M. Maggs, A. Meyerson, J. Saks, and R. K. Sitaraman, "Algorithms for constructing overlay networks for live streaming," *CoRR*, vol. abs/1109.4114, 2011.
- [15] F. Zhou, J. Liu, G. Simon, and R. Boutaba, "Joint optimization for the delivery of multiple video channels in telco-cdns," *IEEE transactions on network and service management*, vol. 12, no. 1, pp. 87–100, 2015.
- [16] A. Fressancourt, C. Pelsser, and M. Gagnaire, "Kumori: Steering Cloud traffic at IXPs to improve resiliency," in *Proc. of DRCN*, 2016.
- [17] M. Mathis, J. Semke, J. Mahdavi, and T. Ott, "The macroscopic behavior of the tcp congestion avoidance algorithm," *ACM SIGCOMM CCR*, vol. 27, no. 3, pp. 67–82, 1997.
- [18] J. Desrosiers, F. Soumis, and M. Desrochers, "Routing with time windows by column generation," *Networks*, vol. 14, no. 4, pp. 545–565, 1984.
- [19] Y. Xiao, K. Thulasiraman, and G. Xue, "GEN-LARAC: A generalized approach to the constrained shortest path problem under multiple additive constraints," in *Proc. of ISAAC*, 2005, pp. 92–105.